

Debbie Probst Day Care Center. Numerous others participate in the protocol studies.

The Division of Oncology also firmly believes that excellence in patient care and in teaching programs is best achieved where there is a continuing pursuit of new knowledge. Each of the six full-time faculty members in the Division is actively engaged in cancer research. The clinical research efforts are concerned with the refinement and development of more effective methods of treatment. New chemotherapy is being sought and tested. Better combinations of chemotherapy, and of chemotherapy with other methods (surgery, radiation, immunotherapy), are also being developed.

#### Debbie Probst Oncology Day Care Center

The Division's new, modern, outpatient clinic was designed in response to the physical and emotional needs of cancer patients undergoing chemotherapy. Located on the lower level of the Stanford Hospital, it is designed as a self-contained unit, convenient and comfortable for both patients and attending medical personnel.

Three kinds of treatment rooms are provided, including some for observation or for lengthy (six to eight hours) infusions that formerly had required hospitalization. Efficient service to patients is facilitated by a television monitoring system (see discussion of Motorola system below), a computer-based medical record system (see discussion of TOD below), and facilities for preparing chemotherapy, analyzing blood, and viewing x-rays.

#### Information Display System

When the Oncology Day Care Center was designed, plans were made for an automated scheduling and information display system. This system was developed in conjunction with the Motorola company and is now in operation in the clinic. The microprocessor-based system signals alphanumeric video information to remote locations via video cables. Scheduling secretaries keep appointment records on an associated floppy disc, and on any given day four video display monitors in the oncology conference room are used to display the day's schedule, relevant lab test results for the outpatients being seen that day, room assignments, and the name of the oncologist who will be attending each patient. At present all data are entered by secretarial personnel and there is no hands-on interaction between the physicians and the small computer.

Time-Oriented Databank (TOD)

For the last several years the Division of Oncology has also been using the time-oriented record keeping system (TOD) originally designed by Dr. J. Fries for use in the Stanford Immunology Clinic [25],[58]. The data and all TOD programs are stored in the Stanford campus computer facility, an IBM 370/168. The emphasis in the design of the TOD system has been the analysis of large amounts of data on a body of similar patients, not on interactive record keeping in the clinical setting itself. Thus there are large amounts of data on Stanford oncology patients, stored by dates of clinic visits, kept on this distant computer for retrospective analysis. TOD provides several programs for statistical analysis of correlations, assessing prognosis by attribute matching, and assisting with other tasks that have traditionally required arduous chart review. Since the data are not currently being used for the care of individual patients, there may be a time lag of weeks before transcriptionists extract the relevant data from paper-based oncology outpatient charts and enter them into the TOD databank.

Oncology Treatment Protocols

As mentioned above, the Division of Oncology is active in clinical research and has many patients being treated under research protocols. There are currently about 30 operational protocols, about half of which are active in the sense that several patients are enrolled in the treatment plan at any given time. Many of the protocols are designed and overseen by Stanford oncologists, but there are also cooperative studies involving Stanford and several other institutions. In many cases, the cooperative studies are overseen by the Northern California Oncology Group (NCOG) which has its headquarters very near the Stanford campus. Each protocol is described by a lengthy article, often 45-60 pages, that explains the justification for the therapeutic approach, outlines criteria for patient selection for the study, describes therapeutic options, and details the specific chemotherapy doses, dose modification, and laboratory and clinical data that must be obtained on each visit. It is quite impossible for any single individual to know the details of all 30 protocols. This is a particularly great problem because the physicians seeing oncology outpatients include fellows, residents, and medical students; these individuals have limited oncology experience and, in the case of house staff and students, generally rotate through oncology for only 4-8 weeks at a time. (See [41] for discussion of one approach which emphasizes use by primary care physicians, but has not emphasized a well-designed human interface.)

II.A.3. Rationale

The rationale for the proposed research has largely been described in previous sections. In short, there has been limited success of statistical, data retrieval, and decision analysis programs in dealing with the judgmental knowledge of expert physicians and the uncertainties of medical data. We have made encouraging strides in the development of symbolic reasoning techniques for application to clinical decision making and believe that the time is now appropriate for the clinical implementation of such a system. Only then will it be possible to assess the power of capabilities which have been designed to make consultation systems acceptable to physicians. Although we recognize that the short term impact of such systems is limited by the current state of the art in computer science, the impetus for appropriate basic research and development of new interactive techniques will come largely through the lessons learned in undertaking clinical implementations. Since techniques already exist that have potential for considerable short-term clinical impact, we believe it is now appropriate to spend part of our time on a project for clinical use.

Although our interest is in the development of systems for offering any kind of subspecialty expertise to primary care practitioners, the initial application selected has been the management of complex therapy protocol information in an outpatient oncology clinic. This domain was selected for a number of reasons:

(1) There are large amounts of information in the protocols but relatively little inferential complexity; those problems that have prevented us from attempting clinical implementation of the MYCIN System for infectious diseases can therefore largely be avoided.

(2) There is a small core of faculty members and oncology fellows who are largely responsible for patients in the day care center. Hence a relatively small number of individuals will need to be introduced to the consultation system, and their continuing roles in the clinic will heighten their chances of becoming comfortable with computer-based techniques.

(3) There is already an awareness of, and involvement with, computers in the Oncology Day Care Center (in the form of the information display system previously described and associated video display monitors). Thus, although there is not yet hands-on computer use by oncologists in the clinic, computer-related hardware is evident and accepted by the clinicians at the outset

of the proposed research. Many fellows and faculty also use the TOD system for clinical research and thus have limited, but very positive, experience with computer use.

(4) Although the application of symbolic reasoning techniques to the protocol management problem will not tax many of the capabilities we have developed in the MYCIN context, it is precisely this simplicity which makes the problem appealing as a first clinical venture. If the information handling task can be implemented relatively easily within the EMYCIN formalism, as we expect it can, then we will be able to concentrate initially on issues of making the system's reasoning and knowledge base understandable as well as making the system's interaction acceptable to physicians.

(5) The initial investment in establishing a role for interactive computing in the oncology outpatient setting at Stanford will have considerable potential for facilitating interactions between our protocol management system and the Division of Oncology's current computer-related efforts (the information display system, and the time-oriented databank). We envision some challenging extensions to the consultation program whereby physicians interacting with the protocol management system may simultaneously benefit from direct connections between our computer and the other oncology systems.

II.B. Specific Aims

We propose core research as well as new demonstrations of the clinical usefulness of present capabilities developed under MYCIN research.

As has been discussed, we have identified an important clinical problem in the outpatient oncology clinic at Stanford, and have begun a collaboration with members of the oncology division to develop and implement a Protocol Management System (PMS) for use in the oncology clinic. Our proposal is to demonstrate that computer-based reasoning and interactive techniques developed during MYCIN research can be effectively applied to an important clinical problem, namely the management of oncology protocol data.

The infectious disease domain with which we have been involved to date involves complex reasoning and computing problems that we feel prevent the short term development of a clinically useful infectious disease consultation system. The oncology problem, on the other hand, involves large amounts of knowledge but rather simple reasoning that current techniques should be able to manage effectively. The complexities of infectious diseases, however, have provided a particularly appropriate domain for devising new computing approaches while analyzing clinical reasoning. These difficult problems remain major research interests of our group. We propose spending approximately half our time continuing to work on basic tools for expert medical consultation systems, using the current content of the infectious disease knowledge base without any efforts to extend its scope in the short term.

Specifically, our aims during the five years of proposed research are:

Artificial Intelligence Objectives

(1) To implement and evaluate recently developed techniques designed to make computer technology more natural and acceptable to physicians;

(2) To extend the methods of rule-based consultation systems to interact with a large database of clinical information;

(3) To continue basic research into the following problem areas: mechanisms for handling time relationships, techniques for quantifying uncertainty and interfacing such measures with a production rule methodology, approaches to acquiring knowledge interactively from clinical experts. These are some of the problems we have identified that have prevented the MYCIN infectious disease application from being clinically implemented as yet.

### Oncology Clinic Objectives

We plan to develop and implement a Protocol Management System (PMS), for use in the oncology day care center, with the following capabilities:

(1) To assist with identification of current protocols that may apply to a given patient;

(2) To assist with determining a patient's eligibility for a given protocol;

(3) To provide detailed information on protocols in response to questions from clinic personnel;

(4) To assist with chemotherapy dose selection and attenuation for a given patient;

(5) To provide reminders, at appropriate intervals, of follow-up tests and films required by the protocol in which a given patient is enrolled;

(6) To reason about managing current patients in light of stored data from previous visits of (a) the individual patients (b) the aggregate of all "similar" patients.

### Advantages over present paper-based protocol files:

(1) Can be kept readily accessible and up-to-date;

(2) Can provide customized patient-specific calculations and advice not possible with a manual system;

(3) May be augmented to provide important additional capabilities once interfaced with a patient data base (e.g., the time-oriented data bank [TOD] already used for retrospective data analysis by the oncology division);

(4) Can provide customized explanations of protocol information and the specific recommendations made by the management system;

(5) Can improve the quality of clinical research by encouraging enrollment of all patients in an appropriate protocol, and assuring that necessary data are obtained to assure uniformity of information on patients in the individual study groups;

(6) Can improve the quality of patient care by:

(6a) Saving time by making protocol information easily available, thus decreasing the waiting time patients must now occasionally sustain while physicians track down necessary protocol information;

(6b) Making certain that important tests are done to screen for potentially serious toxicity of the powerful agents used in cancer chemotherapy.

II.C. MethodsII.C.1. Overview

Our general approach to the research will be to emulate the organizational and technical framework used during development of several interdisciplinary computing efforts involving Stanford's Heuristic Programming Project (HPP), of which Prof. Buchanan is co-director. The cohesiveness of project workers has always been facilitated by a weekly group meeting in addition to smaller working sessions at other times. At group meetings both computer science and clinical personnel have opportunities to present their work and give and receive suggestions regarding further efforts. We believe it is important that the physicians and computer scientists get to know each other and their motivations for involvement in the project very well. For example, the computer scientists working on MYCIN have all learned a great deal about infectious diseases, and some have even taken formal courses in microbiology at the medical school. Similarly, the clinicians have been encouraged to understand the program in depth and even to try some programming. We would expect similar relationships to develop among the computer scientists and oncologists working on the proposed research. Only in this way can both computer science and clinical concerns be taken adequately into account during system design and implementation.

In addition to the development of the PMS for the oncology clinic, we anticipate continued research into the basic science issues discussed previously. As has been noted, we have already identified several problems that must be solved before complex reasoning programs such as MYCIN can be made available for clinical use. We also anticipate that work in the oncology domain will uncover new problems, not previously encountered, that may require significant modification or redesign of the EMYCIN formalism. Thus we envision two parallel but highly interrelated efforts:

(1) development of the PMS for the oncology clinic, using EMYCIN and writing new production rules to embody the protocol knowledge that will be needed for consultation sessions;

(2) continued mapping of basic science research, from the core research section of this program, into the oncology problem domain in order to facilitate complex decision making and acceptable consultations in the clinical setting.



II.C.2. Oncology Protocol Management System

The first year of research on the PMS will be spent developing the program before it is made available in the clinic. Years 2-3 will be devoted to revisions and extensions of the protocol management system in light of initial experience with a knowledge base about oncology. Years 4-5 will be devoted to revisions and extensions of the basic methodology, as well as of the working system, to facilitate use of a clinical data base for patient management in oncology and related disciplines. We expect that the five years will be spent as follows:

(1) We will begin by selecting the 2 or 3 most frequently used oncology protocols (e.g., oat cell carcinoma of the lung, Hodgkin's Disease, non-Hodgkin's lymphoma). The extensive knowledge in these documents will be extracted by the oncologists working closely with those who know the EMYCIN formalism well. Although much of the knowledge can be represented in typical EMYCIN production rules, we anticipate that some of the information may be best contained in alternate representation schemes. We therefore expect that new techniques for interfacing EMYCIN production rules with tabular data or algorithmic structures may be necessary. Most problems that will arise along these lines should develop during codification of the first few protocols; since the protocols all follow a similar structured format, it is unlikely that new problems will arise when the 29th or 30th protocol is being considered.

(2) EMYCIN's knowledge acquisition capabilities remain somewhat rudimentary (see next section), so we expect that most new rules will be explicitly written by members of the research group.

(3) Specific attention will be given to extracting knowledge regarding patient eligibility for a protocol, tests and films needed at various stages of treatment, therapeutic alternatives available, and patient-specific indications for modifying or withholding therapy. We recognize that these are the protocol details that are often most difficult for the oncologists to remember or to extract easily from a lengthy written protocol (an up-to-date copy of which may not even be readily available in the clinic).

(4) Once the knowledge has been codified, we will begin internal testing by interfacing the new production rules and knowledge structures with the EMYCIN program. Of particular interest will be the adequacy of EMYCIN's explanation capabilities when interfaced with this new knowledge base.

(5) Modifications will be made to the EMYCIN system in response to suggestions made by the clinicians working on the project as they gain experience with its capabilities. Of primary concern will be an assurance that the human interface is sufficiently comfortable that the other Division oncologists will be willing to experiment with the system once it is introduced in the clinic.

(6) After these first few protocols are operationally managed by the PMS as described, the system will be introduced in the Oncology Day Care Center. Orientation sessions will be given to the clinic oncologists, and suggestions for further refinements solicited.

(7) The next 3-5 therapy protocols will then be added to the system, with appropriate notification to clinic physicians when a new protocol is available for PMS access.

(8) Based on the experience gathered in codifying the first several protocols, a protocol-entry system with editor will be developed. This should greatly facilitate the entry of the remaining protocols, which we anticipate should be fully codified by the end of year 2.

(9) Anticipating an interface with the TOD system described earlier, plus progress in the basic research that we will be undertaking simultaneously, we will next begin to store patient-related data in TOD format within the PMS. Much of the information in the TOD Databank is also required by the PMS, so there would be minimal if any additional effort required of the PMS user.

(10) Assuming a breakthrough in the representation and management of time-dependent variables, we would anticipate that the PMS capabilities would be greatly augmented by access to patient data stored in TOD format. During Years 4-5 we would attempt to begin the implementation of this kind of interface between TOD and the PMS.

All research described above would occur on a research computer that could not guarantee reliable service to the oncology clinic. We therefore recognize that we cannot initially undertake any tasks crucial to clinic or Division operation. The clinic must be able to continue to function even when our tool is unavailable for scheduling or hardware reasons.

Therefore, when the PMS is ready to progress into a more integral role in clinic operation, we would anticipate, in a separate proposal, the need for a dedicated machine to permit reliable clinic service. We recognize that many of the most interesting and challenging decision making tasks, including those related to the use of symbolic reasoning techniques in conjunction with large databases, can not be made available to clinicians without a dedicated computer, but that this is beyond the scope of the present proposal.

## Project 2

### III. A WORKBENCH FOR KNOWLEDGE REPRESENTATION

#### III.A. Objectives of the Research and their Significance

Our primary strategy for conducting our investigations has been to allow the problem to condition the choice of scientific paths to be explored. Projects One and Three, dealing with problems in oncology outpatient consultations and with the clinical laboratory, are the newest examples.

We are also motivated, however, by the importance (to us and others) of generalizing our techniques and systematizing our methodology. This is a normal part of the activity of cumulating the results in our science, in which the experiments we choose to generalize upon are the experimental systems we construct for different domains of knowledge. In Computer Science, one effective method of cumulating our growing understanding is construct software packages that are the working manifestation of what we believe we have come to understand. These packages allow us to transfer yesterday's "experimental technique" into tomorrow's "tool" for accelerating the research.

These packages also allow investigators in other institutions to build rather directly upon the results of our work, thereby amplifying the science as a whole. It is particularly appropriate to cumulate our knowledge as software packages in the SUMEX-AIM community in which the users share the same computer and system.

We have sought to extract from our various projects the uniformities that have general applicability; to eliminate the ad-hoc features that accrue in any large-scale programming effort; and to build helpful "front-end" interfaces that will allow others to couple smoothly to our work.

A number of such packages are beginning to emerge. We propose to continue their development and test; and to merge them appropriately into a larger software system that (for lack of a better term) we refer to as the "knowledge representation workbench".

The Stanford group is fortunate to have the collaboration of the Missouri group to act as a test-and-evaluation site for this workbench concept. It is expected that much of the research of Project Three will be done using the emerging "workbench".

We propose the following major objectives:

1. To develop AI technology as software packages that solve general classes of problems.
2. To actively disseminate the technology by publication and by encouraging pilot projects using the technology.
3. To apply these packages to medical applications forming collaborations over time as opportunities arise.

### III.B. Background and rationale

Artificial intelligence research at the Heuristic Programming Project has concentrated on programs having real-world applications. Each program has been a case study for representing and manipulating the task-specific knowledge for an application. Feigenbaum [22] has described this case study approach as essential in building a science for "knowledge engineering".

Because the cases have been carefully chosen, the experience from this approach has accumulated. For example, the GAL program [53] was developed recently for inferring DNA structures from enzyme digest data. This program used the Generate-and-Test paradigm — in which the combinatoric output of a complete and canonical generator of possible structures is limited by pruning rules which use the digest data. That basic approach was pioneered by the DENDRAL [11] program ten years ago. With DENDRAL as an example, the development of this analogous program was completed in only two months.

This example shows how the accumulation of theory speeds the development of new AI programs. Significantly, the Heuristic Programming Project has also accumulated methods — in the form of software packages which can perform specific symbolic computations. These packages are the state-of-the-art tools for applied artificial intelligence. A trained "knowledge engineer" can combine these packages to create computer programs for new applications — without having to re-program the solution of standardized subproblems which have been solved before.

EMYCIN<sup>1</sup> is an example of such a package. It is the domain independent core of the MYCIN [51] program for the diagnosis of infectious diseases. EMYCIN provides a framework for building consultation programs in various domains. It uses a production rule mechanism and backward-chaining control structure during the solution phase and has dialogue facilities for acquiring a production rule knowledge base. An example of an application of EMYCIN is the PUFF system for diagnosing pulmonary function disorder. PUFF was the product of a collaboration with the Pacific Medical Center in San Francisco. the first version of PUFF was built in the following way. One hundred cases, carefully chosen to span the variety of disease states were used to extract 55 rules. The knowledge base was created with EMYCIN and then tested with 150 additional cases. Agreement between PUFF and the human expert was excellent and a later version of PUFF is now in routine use at PMC. The first version of PUFF was created in less than 50 hours of interaction with experts at PMC and with less than 10 man-weeks of effort by the knowledge engineers. Other applications of EMYCIN will be discussed in the Section III.C..

The example shows that methods, in the form of usable computer packages, have now been developed. These packages reflect the commonalities we now perceive among separate applications. They are the recently available tools of applied artificial intelligence — programs providing practical symbolic methods for common problems.

Our current repertoire of "methods" packages also include the Unit Package, and AGE-1. The EMYCIN program, as discussed above, is based on production rule technology and has been successfully applied to diagnosing pulmonary function disorders and consulting on structural analysis in an engineering application. The Unit Package [52] is based on the so-called "frames" approach and is being applied to experiment planning in molecular genetics. The AGE-1 program is based on the HEARSAY [20] "cooperating knowledge sources" model and is the product of experience with the SU/X and SU/P [43] programs.

New applications are currently being developed for each of these packages. Heiser and Brooks at the University of California at Irvine are using EMYCIN to develop a psychopharmacology consultant, termed HEADMED [34]. Blum [5] has proposed using the Unit Package in a system which will combine statistical methods and artificial intelligence techniques to perform studies on a clinical database. Several other applications have been proposed and are under consideration.

---

<sup>1</sup>The name "EMYCIN" comes from "essential MYCIN", the MYCIN reasoning framework without any domain-specific knowledge.

We propose to continue the development and application of these packages and to develop new ones as results become available from core research.

### III.B.1. Relating the Workbench to Core Research

Over the five year course of this research, there will be a movement of topics from core research into developed packages for the workbench. Our overall strategy has two main thrusts:

1. To expand the problem solving capabilities of the workbench by developing more sophisticated methods of symbolic reasoning.
2. To expand the capabilities of existing packages following core research in other topics — knowledge acquisition, knowledge integration, tutoring, and explanation.

This mode of research reflects a bias towards the creation of systems to perform specific tasks. First an approach to problem solving is developed and tested in a task domain. Then research in other topics follows. Three methods of problem solving are discussed in this proposal and elaborated in the following. The simplest of these is a backwards chaining approach — exemplified in EMYCIN — which links together the premises and conclusions of rules to construct a direct line of reasoning. The next level of sophistication in these packages is represented in the AGE-1 which is based on the HEARSAYII [20] architecture. AGE-1 allows (1) both data-driven and goal-driven reasoning and (2) reasoning at different levels of abstraction. This architecture has been used effectively by Stanford researchers in a signal-processing application [43]. Providing other AI capabilities — such as explanation or knowledge acquisition — is more difficult in AGE-1 than in EMYCIN. The next level of sophistication appears in a proposed "planning package" which is expected to grow out of on-going research in the MOLGEN project. This approach to planning formalizes the selection of what to do next as a choice in any of several problem-solving "spaces". The viability of the latter problem-solving method is still being tested and essentially none of the other system capabilities have been developed.

The following is a list of several AI issues discussed in this proposal. These will be explored within some formalisms

already developed by us, EMYCIN, AGE-1, and the Unit Package <sup>2</sup>, as well as new formalisms, e.g., the Planning Package as the need arises. The planning package is expected to materialize at the end of some core research which is currently in progress.

Problem Solving  
Knowledge Acquisition  
Explanation  
Tutoring  
Knowledge Compiling  
Time-Dependence  
Meta-Knowledge

### III.C. Methods of procedure

This section describes our plan for creating an integrated collection of well-designed software packages, which can be combined by a knowledge engineer to meet the needs of a specific application. In this section we will show examples of each of the packages and discuss the nature of their applications. We will also discuss the work proposed for further developing the packages.

There is a great deal of overlap in the proposed work among the packages. While the packages reflect different approaches to problem solving and differ in their state of development, analogous lines of research are proposed in each. The EMYCIN package, which is the most developed, uses the the simplest approach to problem solving and has the broadest range of proposed work following several lines of core research. As discussed already in Section III.B.1., similar lines of development are planned later in the grant period for the other packages.

#### III.C.1. EMYCIN

The EMYCIN ("Essential MYCIN") project is an attempt to provide a framework for building consultation programs in various domains. It uses the domain-independent components of the MYCIN

---

<sup>2</sup>The Unit Package is a passive representation package and does not provide any software for problem-solving. It is being used, however, as a representation medium for the Planning Package and can also be used in conjunction with AGE-1.



system, notably the production rule mechanism and backward-chaining control structure. Then for each particular consultation domain, the system builder supplies the rules and parameters of that domain to produce a functioning program. Work on the EMYCIN project is devoted to providing a useful environment for the new system builder, with emphasis on speeding the acquisition and debugging of the knowledge of the new domain.

#### III.C.1.a. An Example of EMYCIN — The PUFF Application

The PUFF system for the interpretation of laboratory measurements from the pulmonary function laboratory. The EMYCIN system was used as base upon which 60 production rules concerning the presence of pulmonary disease were created. The data from over 100 cases were used to create the rules by the pulmonary physiologist in cooperation with the biomedical engineers who instrumented the laboratory and Stanford computer scientists who had previous experience with the MYCIN program.

Figure 1 shows several rules created during the development of the system. These rules are used to create a complete report including the input measurements, historical information, and the measurement interpretation. Figure 2 shows a copy of this report.

```

IF  0 < DLCO < 80      (DLCO is the measurement of diffusion
                        capacity for Carbon Monoxide)
THEN "Low diffusing capacity indicates loss of alveolar
capillary surface which is "
    IF  70 <= DLCO < 80 THEN "mild"
    IF  60 <= DLCO < 70 THEN "moderate"
    IF  0 <= DLCO < 60 THEN "severe"

IF The severity of obstructive airways disease of the patient is
   greater than or equal to mild, and
   The degree of diffusion defect of the patient is greater than
   or equal to mild, and
   The total lung capacity measured by the body box (TLCB) is
   greater than 110 percent of predicted,
THEN "The low diffusing capacity, in combination with
obstruction and a high Total Lung Capacity, would be consistent
with a diagnosis of emphysema."
    The subtype of obstructive airways disease is emphysema.
```

Figure 1. Typical PUFF interpretation rules. Conclusions are made for internal system use and for inclusion in the summary.

PRESBYTERIAN HOSPITAL OF PMC  
CLAY AND BUCHANAN, BOX 7999  
SAN FRANCISCO, CA. 94120  
PULMONARY FUNCTION LAB

DOE JANE 5S2  
P336666.  
DR. SMITH, JOHN

WT 56.7 KG, HT 166 CM, AGE 58 SEX F  
SMOKING 40 PK YRS, CIG 1.0 PK QUIT 0, PIPE 0 QUIT 0,  
CIGAR 0 QUIT 0  
DYSPNEA-W/MILD-MOD. EXER, COUGH-NO, SPUTUM-LT 1 TBS, MEDS-YES  
REFERRAL DX-CORONARY ARTERY DISEASE, PRE OP

\*\*\*\*\*TEST DATE 10-26-78

		PREDICTED (+/-SD)	OBSER(%PRED)	POST DILATION OBSER(%PRED)
INSPIR VITAL CAP (IVC)	L	3.1(0.4)	3.0 (98)	
RESIDUAL VOL (RV)	L	2.1(0.3)	3.0 (140)	3.5 (166)
TOTAL LUNG CAP (TLC)	L	5.2(0.7)	6.0 (116)	6.5 (125)
RV/TLC	%	40.	49.	53.
FORCED EXPIR VOL (FEV1)	L	2.6(0.3)	2.1 (81)	2.1 (84)
FORCED VITAL CAP (FVC)	L	3.1(0.4)	2.9 (96)	3.0 (98)
FEV1/FVC	%	83.	70.	71.
FORCED EXP FLOW 200-1200L/S		4.2(0.8)	4.5	4.4
FORCED EXP FLOW 25-75% L/S		2.9(0.7)	1.5	1.5
FORCED INS FLOW 200-1200L/S		2.9(0.6)	2.9	2.9
AIRWAY RESIST(RAW) (TLC= 6.0)		1.1(0.5)	1.6 (HIGH)	1.4
DF CAP-HGB=14.4 (TLC= 5.3)		25.	17.2 (68)	(69%IF TLC= 5.2)

\*\*\*\*\*

INTERPRETATION: Elevated lung volumes indicate overinflation. In addition, the RV/TLC ratio is increased, suggesting a mild degree of air trapping. Forced vital capacity is normal but the FEV1/FVC ratio is reduced, suggesting airway obstruction of a mild degree. Reduced mid-expiratory flow indicates mild airway obstruction. Obstruction is indicated by curvature in the flow-volume loop of a small degree. Following bronchodilation, the expired flow shows slight improvement. This is confirmed by the lack of change in airway resistance. The low diffusing capacity indicates a loss of alveolar capillary surface, which is moderate.

CONCLUSIONS: The low diffusing capacity, in combination with obstruction and a high total lung capacity would be consistent with a diagnosis of emphysema. The patient's airway obstruction may be caused by smoking. Discontinuation of smoking should help relieve the symptoms.

#### PULMONARY FUNCTION DIAGNOSIS:

1. Mild Obstructive Airways Disease.  
Emphysematous type.

Robert Fallat, M.D.

Figure 2. Sample PUFF Report

III.C.1.b. Applications of EMYCIN

To date, EMYCIN has been successfully applied at Stanford to the domains of pulmonary function (PUFF) [37] and structural analysis (SACON) [3].<sup>3</sup> EMYCIN is also being applied to clinical psychopharmacology [34] at the University of California at Irvine.

III.C.1.c. Proposed Work for EMYCINSYSTEM-BUILDING TOOLS

## 1) Acquisition of Knowledge — Acquire the

---

<sup>3</sup>SACON (Structural Analysis Consultation): The purpose of the consultation is to provide advice to a structural engineer regarding the use of a structural analysis program called MARC. The MARC program uses finite-element analysis techniques to simulate the mechanical behavior of objects. The engineer typically knows what he wants the MARC program to do, e.g. examine the behavior of a specific structure under expected loading conditions, but does not know how the simulation program should be set up to do it. The MARC program offers a large (and, to the novice, bewildering) choice of analysis methods, material properties, and geometries that may be used to model the structure of interest. The user must learn to select from these options an appropriate subset that will simulate the correct physical behavior, preserve the desired accuracy, and minimize the (typically large) computational cost. The goal of the SACON program is to bridge this gap, by recommending an analysis strategy. This advice can then be used to direct the MARC user in the choice of specific input data, e.g. numerical methods and material properties.

The performance of the SACON program matches that of a human consultant for the limited domain of structural analysis problems that was initially selected. To bring the SACON program to its present level of performance, about two man-months of the expert's time were required to explicate his task as a consultant and formulate the knowledge base, and about the same amount of time implementing and testing the rules (this estimate does not include the necessary time devoted to meetings, problem formulation, demonstrations and report writing).

framework, vocabulary, and decision rules of the domain from the expert.

2) Rule Checking — Check syntax and semantics of new rules and check for possible conflict with existing rules.

3) Alternative Models for Reasoning under Uncertainty — Provide the system builder with a fixed set of alternative methods for propagating degrees of certainty in the reasoning chains.

4) Time-Dependent Features — Enable the system to make use of parameters whose values change with time.

5) Meta Knowledge — Add capabilities for using meta-rules and other meta-level knowledge.

In addition, we propose extending the power and flexibility of the present system in the following ways:

#### DOMAIN-INDEPENDENT CONSULTATION SYSTEM

1) Answering Questions — Incorporate question-answering capabilities into the system.

2) Tutoring — Couple the system to a tutoring program to teach the contents of the knowledge base.

Many of these items involve substantial research before we understand the best way to add them to the program or even what, precisely, needs to be added. We present below our best ideas on the approach we will take, but wish to emphasize that the nature of the solution may change as our research progresses.

The products of the research will be presented in scientific papers and in an integrated computer program that can be used by scientists to encode their own knowledge of their domains for reasoning about difficult problems.

III.C.1.d. Acquisition of Knowledge

The preliminary facilities for acquiring knowledge (called TEIRESIAS [Davis76]) developed in the context of the MYCIN application will be incorporated into EMYCIN for use by experts when building any consultation system. This facility will allow an expert to specify the major parameters of a consultation. Then, following a consultation, the system will show the expert the values of these parameters, and ask for verification that they are correct. If the values are not correct, the system will explain to the expert the line of reasoning that led to the incorrect values. This allows the expert to pinpoint an error in the system's rule set, which the expert can then repair by adding, deleting, or modifying rules.

In addition to incorporating the existing rule-acquisition facility, we plan to automate the acquisition of a large portion of the initial knowledge that is required in building a consultation system. The system will prompt an expert through an intermediary for the conceptual framework, vocabulary, and major lines of reasoning of the domain before any rules are entered. The conceptual framework includes the definition and hierarchy of objects or states that will be used to structure the reasoning process (called the "context tree") as well as the attributes and values of these objects that will be used for writing rules. Numerous internal pointers needed for correct associations among concepts will be set up automatically at this time.

Improvements to Teiresias

The TEIRESIAS facility, for interactively debugging the rule base, is most useful when the knowledge base is reasonably well developed and the necessary changes to the rule and parameter base are small. This facility is currently being improved primarily by using the existing question-answering system to explain the system's lines of reasoning [48], and by using a new English parser based on a semantic grammar to understand any rule additions or changes from the expert [8].

An EMYCIN sketchpad As a result of our recent experience eliciting a rule base for structural mechanics [3], we have found it useful to characterize the knowledge acquisition process as occurring in a number of distinct phases.

The first phase corresponds to making initial decisions about the typical advice the consultant will give and the major reasoning steps the consultant will use.

This is followed by an extended period of defining parameters and objects and then, using this initial domain vocabulary, developing a substantial portion of the rule base. This process, lasting approximately 2 months in the structural analysis case, captures enough domain expertise to allow the consultation system to give advice on the large number of common cases.

In the final phase, further interactions with the expert tend to refine and adjust the established rule base, primarily to handle more obscure or complicated cases.

Future research on knowledge acquisition will explore the design and implementation of interactive facilities to be used during the early phases of the knowledge base design. In particular, methods will be developed for rapidly acquiring and manipulating definitions of the context tree of objects, their major parameters, as well as the major problem solving strategies to be used by the consultant.

During the initial passes at defining objects, the system would begin to acquire some detail about the actual methods (the rule sets) that will be used to reason about the major parameters of the consultation. For each of these parameters the expert typically knows what major factors and subgoals will be relevant to concluding the parameter. These factors can be specified by the expert, but need not be acquired in detail until the system actually must begin gathering the rules for determining these important parameters. In this manner, the expert can be free to concentrate on the more general aspects of the problem solving process without having to be bothered with the specification of detail.

Using the EMYCIN sketchpad, the expert and intermediary would develop and acquire substantial portions of the knowledge base and an explicit representation of the overall reasoning strategy that the program will use to advise about the user's problem. This framework and knowledge of overall strategy can be used later to motivate explanations of the system's lines of reasoning produced by the question-answering system. We intend to investigate ways that this knowledge about the major parameters could be used by TEIRESIAS (during the later phases of the knowledge acquisition process) to explain how and why a particular, incorrect conclusion was made.

#### Rule Checking

While the production rule format permits any executable LISP expression as the premise or action of a rule, not all LISP forms make reasonable rules. Common syntactic errors include misspellings, misplaced arguments, parenthesis errors and incorrect classification of the rule; such errors generally result from inaccurately inputting the rule, and if left undetected, may cause the rule to fail, or even cause runtime errors. Semantic errors can result if a new rule is inconsistent with existing rules, or is incomplete, failing to take into account all the factors necessary for the conclusion.

We plan to do extensive checking of each new rule entered into the system. We hope thereby to catch most errors at rule entry time, rather than finding them during later consultation runs when it is harder to (a) isolate the effects of a faulty rule and (b) correct any problems which result.

Syntactic checking is fairly straightforward. The rule checker needs to know about the syntax of each argument to the predicates which make up a rule. This knowledge exists in the form of predicate templates, which have long been used by other parts of the system to "read" rules. The rule checker's use for them is, in effect, to make sure the rules are "readable". For example, the template for the predicate SAME is

(SAME CNTXT PARM VALUE),  
for which a typical instance from the infectious disease domain might be

(SAME CNTXT IDENT E.COLI).

The rule checker knows from this that a call to SAME should have three arguments: the first must be a legal "context atom", i.e., a variable used to select a binding in the context tree, the second must be a parameter, and the third must be a legal value for that parameter. If any of these is incorrect, the error is easily detectable, and in many cases correctable. Simple spelling errors may be corrected by invoking INTERLISP's spelling corrector, using an appropriate spelling list; e.g., for the PARM slot use the list of all parameters, for the VALUE slot use the list of values legal for the parameter appearing in the PARM slot. Transposed arguments and spurious extra arguments (typically a result of parenthesis errors) are also easily detected by checking against the template.

Another common syntactic error is incorrect classification of a rule, i.e., specification of what type of context it may apply to. In many cases it is possible for a rule checker to completely determine the correct classification, simply by observing which parameters appear in the rule and comparing with the known structure of the context tree. At worst, the checker

could narrow down the possibilities to a small set of nodes of parallel structure.

More subtle errors arise from fundamental "semantic" errors in a new rule, and the processing required to detect such errors is correspondingly more complex. One major type of semantic error is inconsistency of a new rule with existing rules. One rule might subsume another, i.e., one premise is implied by another. For example, with the two rules

A  $\rightarrow$  X

A & B  $\rightarrow$  X,

the first subsumes the second. The error here is that if the second rule succeeds, the first will also, and the information A is contributing twice to the conclusion X. Our certainty factor model is predicated on rule premises being independent; subsumption is a blatant violation of that assumption.

Another possibility is that one rule might contradict another rule or rules. This is trickier. Certainly the two rules

A  $\rightarrow$  X

A  $\rightarrow$   $\sim$ X

contradict each other. But such obvious contradictions are fairly unlikely; more subtle interactions can occur. For example, given a set of rules

A  $\rightarrow$  B, B  $\rightarrow$  C

A  $\rightarrow$  D, D  $\rightarrow$   $\sim$ C

it is difficult to determine whether there is a contradiction except in the special case that all the rules have definite conclusions (CF=1.0). But if the confidence attached to those conclusions is less than definite, there may be no direct contradiction at all, merely conflicting tendencies, perfectly admissible under our certainty factor model. We plan to investigate means of analyzing rules to uncover possible contradictions, measure how great a conflict may exist, and ways to determine if the conflict is a real problem.

Another type of semantic error may occur if a rule fails to take into account all the information relevant to a conclusion. The system can sometimes detect this by means of rule models, which currently consist of statistical observations of the correlation of parameter occurrences in existing rules [15]. These rule models are constructed automatically by reading the rules. As a typical use, if rules mentioning parameter x usually also mention parameter y, then the system might request confirmation of a new rule which considers only x. We plan to increase the richness of the rule model language, to enable better semantic checking of the user's rules, especially during early acquisition phases, when there do not exist sufficient rules to form useful rule models on purely statistical grounds.



For example, the user might wish to describe in some brief fashion the sort of rules he is about to enter, and the system could then make sure the rules are actually consistent with the user's model.

#### III.C.1.e.    Alternative Models for Reasoning under Uncertainty

The method developed for ranking MYCIN's hypotheses based on measures of certainty is an approximate method. It developed from a pragmatic need for measuring the degree of confirmation of a hypothesis based on several non-independent (partially overlapping) pieces of evidence. The certainty factor (CF) model discussed above is a means of combining single "certainty factors" associated with each inference to arrive at a reasonable measure of how strongly the evidence supports each hypothesis.

It is reasonably simple to understand. However, its main drawback lies in the difficulty of associating a CF with a single rule. Because the rules are not independent, the CFs are also not independent. This means that adding a new rule involves looking at similar rules in order to decide how high the CF ought to be set.

For some experts (or problem areas), CFs seem to be more difficult to use than for others. Thus we propose to offer the system builder a choice of evidence accumulation methods. One of them will be the CF scheme already in use. A second will be the likelihood ratio scheme used in the PROSPECTOR system [18], although that requires storing two measures with every inference:  $P[H/E]$  and  $P[H/\bar{E}]$ .

A third method will be a very simple additive measure with thresholding, as proposed by one of the physicians working with MYCIN. In this model, measures of positive and negative evidence are added and subtracted into a total for each hypothesis, with action taken on the hypotheses in the end that lie above the threshold.

Under other funding we are exploring other relationships between evidence and hypotheses. As measures are found that can be fit to new problem areas we will find ways of adding them to the set of available confirmation methods. The important point here is to give the system builder a choice of evidence accumulation schemes, any of which can be used in EMYCIN.